

SAM (REX/Ops) Operations Meeting 2008.01.08

Present: Adam (typist), Randolph, Art, Parag, Angela, Dave, Krzysztof, Andrew

For reports submitted to the meeting...

The report text is in this font

Comments made during the meeting are in this font inline with the report

Releases:

Robert has an 8.4.2 release of the DB server.

Over the holiday week, Randolph and Dehong successfully tested 8.4.1 and retired files with a full project test (including actual file fetching). Uh oh - Krzysztof realizes that the station may have been using the old DB server (v 8.2.0). So if that is the case, then the test needs to be redone. But maybe Dehong did this -- nope, the station has not been restarted since October. Not sure why the test worked, then.

DØ:

Kin Yip to sam-design

Events worth stating :

--- No SAM shifters from ~Dec. 24 (US day time), 2007 to Jan. 3, 2008.

*--- Dec. 24, 2007: Mover D41DLTO and D41ELTO were back in production.
Mover D41CLTO was*

*marked down with volume PRU834L1 stuck in the drive. It'd need further
attention, possibly a service call.*

Now, PRU834L1 is "notallowed" whereas D41CLTO seems in operation.

*--- Jan. 6: (shifter's report) Prague Site showed disconnectivity as
most of the DbServers running there were RED*

*(during second half of the shift). But then they all
came back.*

What is running at Prague? Is it the calibration DB servers?

*-- Jan. 7: (shifter's report) ENSTORE mover:D41ALTO.mover is RED because of tape
ejection error*

Unresolved issues:

#3301: "No batchstatus of samgrid job". Any experts are looking at this ?

*#3298: "phantom batch jobs at grid1.oscer.ou.edu". Parag and Jaime are
working on it.*

Condor people are looking at this.

#3296: "bad MC merge job efficiency". Anybody has taken a look at this problem/complaint ?

Steve and I looked at this -- saw time outs for file transfers - not sure what to do about this yet. Will talk to experts. See ticket for more info.

#3290: "Alert: Warning: D00FPRD1.FNAL.GOV - ORA-error stack (07445[kpopfr()+536])".

Do we want to look at this problem at all ?

Anil thinks this is related to a network glitch. Problem disappeared after a few hours.

#3283: "LCG jobs idle" ? Andrew said that it's problems with LCG RB (resource broker)

There is a limit problem ? Is this ticket similar to #3298 ?

Not sure if this is still an issue for Willem. The resource broker is very slow. LCG may migrate the resource broker to a newer version which is not interface compatible with the old one. That will require work for us. We can still use the old one, but it does not work well. UK site is installing new RB - so our new forwarding node there is a good candidate to try it out. We can still use the old one at Wuppertal for production.

#3275: "Jobs waiting for tarfile". Asking Willem whether this is still a problem as there is no help for him etc.

Willem reports problem resolved (probably UK dCache problems).

CDF:

From Barry:

This is a summary of my SAM shift. Over all things were very quiet.

The cdf-cnaf station has not yet returned after a CNAF power outage on Jan. 4.

samTV has not updated in the last 16 hours and has an obvious change of year problem:

Last File Delivery Wed Dec 31 22:28:20 2008 (34s)

Adam will take a look - could be lingering projects -- today it was more current.

There were some user problems that were attributed to caf restarts. The fundamental problem seemed to be having jobs with 2500 sections. He was told to try 500

This is probably the 1024 connection limit in select. We are looking at options.

Bad Tape IAD846

SSA Response: This tape is marked as being "absent" from the SL8500. We are trying to determine why.

IA1480 also had a bad leader. Leader has been repaired and is awaiting cloning.

On Friday before holidays, dCache timeouts were applied at 5 days total and 3 days inactive.

These changes address clients that can linger forever and potentially cause problems.

One intervention from Angela:

I noticed that ReadOnlyPool r-fcdfdata318-4 had a high number of queued requests, so I moved it

to the TestPools and increased the number of max active movers by 20. I'll keep an eye on the load

and up it again if needed.

- Angela Bellavance

-- Fri Jan 4 11:53:46 comment by...Angela Bellavance --

The pools has drained, so I reset to max number of active movers to normal and moved the pool

back to ReadOnlyPools.

Robert did a similar action during his shift over the holidays

Still one Bad Mover: LTO3_10

K: A person was running a two step job using AC++ - one part using SAM and the other part not using SAM. Did not correctly state in his job that he was not using SAM and got n x 2 consumers. Still corresponding with user.

R: people are running some DFC reports in the DB browser even though no new data have been added.

MINOS:

Arthur Kreymer to sam-design, minos_sam_admin

Minos status 2008 Jan 08

General -

AFS - Timeouts of some servers are continuing

Software upgrade looked promising, but failed. But that failure was probably due to bad hardware. New hardware is being obtained and tests will be repeated.

Enstore -

DCache -

Delayed by E907 (MIP) writing over 12K files to write pools in 4 hours.

Advised to concatenate. Writes were cancelled about a day later.

Would have taken over a week to clear the small files. Were removed by dCache developers. MINOS data taking was halted during this time.

BlueArc/Nexsan -

Major problems 22-26 Dec, including daylong fsck.
Waiting for firmware update for Nexsan array.

SAM -

IT 3128. sam_products for sam_station v6_0_5_22_srm
IT 1894 Need 'sam locate' in C++ API

Oracle -

OPEN ISSUES

IT 1146: cannot test existence of sam tape location
IT 1894: sam locate via C++ api, for simple client needs
IT 1979: Compiler warnings from sam_cpp_api
IT 2071: quadratic project startup time growth
IT 3128. sam_products for sam_station v6_0_5_22_srm

PLAN -

Migration to 9940B tape can proceed anytime

Comments

DB:

D0 :

- On 12/04/07 SAM cut v6_10 on d0ofdev1/dofint1 done

To add composite index on lower(file_name) and retired_date

Drop the unique index on lower(file_name)

Recreated as non-unique on file_name. dropped the unique constraint on file_name
and retired_date

Note: V6_10 will be deployed during downtime.

On 01/03/08 SAM Cut V6_13 on d0ofdev1/d0ofint1/d0ofprd1 - Index on create date for
CACHED_FILE_HISTS, CACHED_FILE_PROJECT_USAGE,
CACHED_FILE_PROJ_USAGE_HISTS tables. done

Note: As per user query time on cached_file_project_usages improved from ~20 to ~0.1s in integration

CDF :

- SAM v6_10 pending on cdfofpr2 - Will be deployed during downtime
- SAM v6_13 pending on cdfofdev/cdfofint/cdfofpr2 No downtime

MINOS:

Alert:

- Buffer Cache Hit Ratio on minosprd is 67%. Maximum is allocated to oracle.
- Need to upgrade the memory on node hosting minosprd. Maurine M. is looking into this.

Details: http://www-css.fnal.gov/dsg/internal/minos/dec13_tuning_minosprd.htm

- SAM v6_10 pending on Minosprd - Will be deployed during downtime
- SAM v6_13 pending on minosdev/minosint/minosprd No downtime

Enstore/dCache:

No info about scheduled downtimes.


Lots of activity over the past week - lots of problems with drives and bad tapes.

See the E9907 issue above in the MINOS section.

There was a MINOS user trying to run over 10,000's of files at once. Complained that this was slow -- Art educated the user.

CDF lent public enstore 40 9940B tapes.

Note Burn rates of stken tapes...



Note large rate of tape usage around October 15th or so.

The rate could have been increased by LQCD or DES. They were very recently moved to LTO3's, so this rate could go back down.

In general, the enstore folks are looking around for blank or recyclable 9940B tapes.

A list of recyclable tapes may be found at...

http://www-stken.fnal.gov/enstore/tape_inventory/RECYCLABLE_VOLUMES

Migration to LTO4 is being considered - especially for D0 raw data (for duplicates, not replacement) - for commissioning of LTO4's.